

Know Your Enemy: Stealth Configuration-Information Gathering in SDN

Mauro Conti
University of Padova
conti@math.unipd.it

Fabio De Gaspari
Sapienza University
degaspari@di.uniroma1.it

Luigi V. Mancini
Sapienza University
mancini@di.uniroma1.it

Abstract—Software Defined Networking (SDN) is a network architecture that aims at providing high flexibility through the separation of the network logic from the forwarding functions. The industry has already widely adopted SDN and researchers thoroughly analyzed its vulnerabilities, proposing solutions to improve its security. However, we believe important security aspects of SDN are still left uninvestigated.

In this paper, we raise the concern of the possibility for an attacker to obtain knowledge about an SDN network. In particular, we introduce a novel attack, named *Know Your Enemy* (KYE), by means of which an attacker can gather vital information about the configuration of the network. This information ranges from the configuration of security tools, such as attack detection thresholds for network scanning, to general network policies like QoS and network virtualization. Additionally, we show that an attacker can perform a KYE attack in a stealthy fashion, i.e., without the risk of being detected. We underline that the vulnerability exploited by the KYE attack is proper of SDN and is not present in legacy networks. To address the KYE attack, we also propose an active defense countermeasure based on network flows obfuscation, which considerably increases the complexity for a successful attack. Our solution offers provable security guarantees that can be tailored to the needs of the specific network under consideration.

Keywords: Network Security, SDN, OpenFlow, Side-Channel Attack, Configuration Information Gathering, Intelligence Gathering

I. INTRODUCTION

Software Defined Networking (SDN) is a network architecture proposed in recent years to address the shortcomings of traditional architectures. SDN posits that the implementation of network functions and the control logic of the network are two separate concepts, and should therefore be separated in different entities. To this end, SDN introduces the concepts of *data plane* and *control plane*: the data plane is comprised of the physical network devices (from here on, called switches) and implements the forwarding functionalities of the

network, while the control plane manages the network logic and decision making process. In SDN, the control plane takes the decisions on how traffic flows are managed and pushes these decisions to the data plane, that will in turn enforce them. This separation between logical control and physical implementation of the network functions provides a high degree of flexibility, which is one of the main reason for the widespread adoption of SDN even amongst big companies [3], [30].

While the programmability of SDN allows for fast prototyping and high adaptability to different scenarios, it also opens new venues for attacks [13]. Indeed, while the decision making process is centralized in the control plane, the enforcement of the decision is distributed throughout all switches, which follow the rules pushed by the control plane. We show that, by exploiting this distributed policy-enforcement mechanism, an attacker can gather intelligence about the control logic of the network in a stealthy fashion. In particular, to run the attack we propose in this paper, the adversary needs to have only a flow table side-channel, i.e., a way of learning which rules are installed, for a single switch (which we call *entry switch*, see Section III). By analyzing the conditions under which a rule is pushed, and the type of such rule, an attacker can infer sensitive information regarding the configuration of the network. The final result is that through a single switch, the attacker can gather information which, in a classical network, would have required access to numerous distinct devices, such as firewalls, intrusion detection/prevention systems, etc. The information gathered can subsequently be exploited to mount different attacks, tailored to the target network, without being detected.

To summarize, our contribution in this paper is as follows:

- We propose a novel, attack, the *Know Your Enemy* (KYE) attack, that allows stealth intelligence gathering about the configuration of a target SDN network. The information that an adversary can obtain ranges from configuration of security tools, such as attack detection thresholds for network scanning, to general network policies like QoS and network virtualization.
- We prove the feasibility and efficacy of the KYE attack through its implementation and a thorough experimental evaluation on a test network.
- We propose a countermeasure to the KYE attack, based on obfuscation of inbound network flows.

II. PRELIMINARIES: OPENFLOW

SDN is a general architectural principle: it broadly defines general guidelines and overall architecture. When discussing SDN in real world scenarios, we refer to specific SDN implementations. In the remainder of this paper we focus on the OpenFlow implementation of SDN, due to its wide adoption in academia [39] and industry (including companies like Google and Microsoft [3], [30]). However, it is worth noting that all our considerations are not specifically tied to OpenFlow, but hold true in general for SDN.

The OpenFlow specification defines the communication protocol between the network controllers in the control plane and the network devices, called OpenFlow switches, in the data plane. Additionally, OpenFlow defines the general architecture and the functionalities of OpenFlow switches. Figure 1 provides an overview of the architecture of an OpenFlow switch. In OpenFlow, each switch maintains a set of flow tables. Each of these flow tables contains a set of flow rules; flow rules are installed on OpenFlow switches by the control plane, and define what actions a switch needs to apply to a certain network flow. A flow rule can match one or more network flows, through the use of wildcards, and a network flow can be matched by one or more flow rules. Through combination of different rules and actions, the controller can instruct the switches to perform arbitrarily complex operations on the packets belonging to a given network flow. One of the main strengths of OpenFlow is the on-demand management of network flows: when a switch receives an inbound network flow for which it has no matching flow rule, it will contact the controller to request the installation of a new rule. This dynamic management of rules allows the controller to evaluate the current state of the network before deciding how to handle a network flow, which allows for the implementation of complex and flexible network policies. The communication between OpenFlow switches and the control plane can be either in plain text, or secured through TLS. The specification does not mandate TLS support, and in many cases commercial OpenFlow switches do not support it [17].

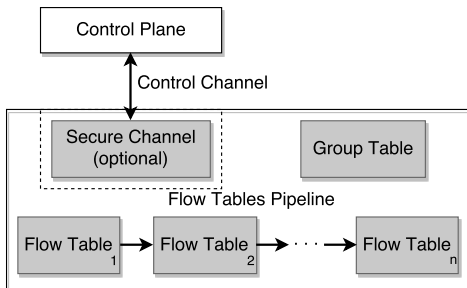


Fig. 1. Overview of the architecture of an OpenFlow switch.

III. ASSUMPTIONS AND SDN ISSUES

While having a logically centralized point of control allows to improve the decision-making process, distributing the policy enforcement introduces new problems with regard to

information disclosure. Where network functions in legacy networks are relegated to the specific devices implementing them, providing higher control over access to their configuration, in SDN they are distributed throughout the OpenFlow switches. Indeed, network policies and functionalities like intrusion detection/prevention systems (IDS/IPS), network virtualization, or access control, are often enforced by the OpenFlow switches, through the application of the flow rules installed by the controller [20], [23], [24]. Unfortunately, this behaviour considerably broadens the attack surface for an attacker.

Indeed, as we show in this paper, by having a flow table side-channel on a single OpenFlow switch (called *entry switch*), an attacker can gather a relevant amount of information regarding the behaviour and configuration of the SDN network the switch belongs to. A flow table side-channel is defined as any mean by which an attacker can learn or infer the content of the flow table of a switch. The KYE attack is independent from how this side-channel is obtained, and the detailed description of how to obtain it is out of the scope of this paper. However, for completeness, we list here some of the possibilities an attacker could leverage:

- Connect to a passive listening port on the entry switch to retrieve the flow table. In fact, most OpenFlow switches can be remotely debugged by means of a passive listening port, which also allows retrieval of the flow table [17], [36]. For instance, an attacker could use the `dpctl` utility [7] on the unprotected listener port of an HP Procurve [1].
- Use Round Trip Time (RTT) variance to infer information about the content of the flow table [21], [32].
- Sniff the control traffic. Indeed, the use of TLS on the control channel is optional [6], and many commercial switches do not support it [17]. Additionally, in most cases OpenFlow switches that support TLS do not implement certificate authentication [36].
- Exploit a backdoor in the OS of the switch to read the flow table or decrypt the control traffic [2].
- Use an hardware device to physically read the flow table [19]. Such devices can acquire the content of the memory of the switch and copy it to an external destination.

It is worth noting that the attacker uses the flow table side-channel only to read the state of the flow table, therefore the overall state of the entry switch is not modified in any way. This means that, even in case of a controller monitoring the integrity of the entire state of the switches, e.g., through direct query [31] or checksum [17], the flow table side-channel would not be detected. Through the KYE attack the attacker can learn all sorts of relevant information, from routing policies, to more complex and important behaviours regarding attack detection and defense mechanisms.

Finally we would like to point out that the KYE attack exploits a structural vulnerability of SDN, which derives from the on-demand management of network flows, that in turn is one of the main features and strengths of this new network

paradigm. Therefore, it is not something that can be easily modified or circumvented.

Threat Model. Our threat model assumes that the attacker has a flow table side-channel for a single OpenFlow switch. In particular, the only abilities of the attacker are (i) sending packets through the target network and (ii) using the side-channel to learn the flow rules that are installed on the entry switch. The attacker can be either internal or external, does not have control over the switch, and does not modify the behaviour of the switch. Furthermore, we assume that the switch only provides the functionality described in the OpenFlow standard [6]. Therefore, the attacker does not modify the software or add capabilities to the switch in any way.

IV. THE KYE ATTACK

From a high level view, the KYE attack follows a general attack strategy to obtain information on a target network. The details of the attack vary based on the specific information the attacker wants to gather (see Section V). However, all instance of the KYE attack share a common kernel, as detailed below. The main idea behind the attack comes from the observation that, in an SDN network, rules are pushed from the controller to the switches only when needed [6]. This behaviour holds true for all types of rules; for instance, the controller will push a flow rule to counter a denial of service (DoS) attack only when one is under way and detected. The KYE attack exploits this on-demand installation of flow rules, allowing an attacker to gather knowledge about which conditions trigger the installation of a given flow rule, as illustrated in Figure 2. The KYE attack is structured in two phases: (A) the *probing phase* and (B) the *inference phase*.

In the probing phase (A), which is repeated numerous times, the attacker attempts to trigger the installation of flow rules on the entry switch (steps 1 through 5 in, Figure 2):

- The attacker sends carefully crafted *probing traffic* through the entry switch in order to trigger the installation of new flow rules. The specific characteristics of the probing traffic depend on what kind of information the attacker is interested in learning (see Section V and Section VI).
- Through the flow table side-channel, the attacker obtains the flow rule (if any) installed in response to the probing traffic.

In the inference phase (B), the attacker analyzes the correlation between the probing traffic generated during the probing phase and the corresponding flow rules installed (step 6 in Figure 2). From this analysis, he can infer what network policy is enforced for specific types of network flows. For instance, if in response to scanning traffic generated in the probing phase the controller installs *drop* rules, the attacker infers that the defense policy against network scanning is traffic filtering. Additionally, by studying the features of the probing traffic, the attacker can potentially infer the trigger conditions for network policies that require a specific trigger before being activated (see Section VI).

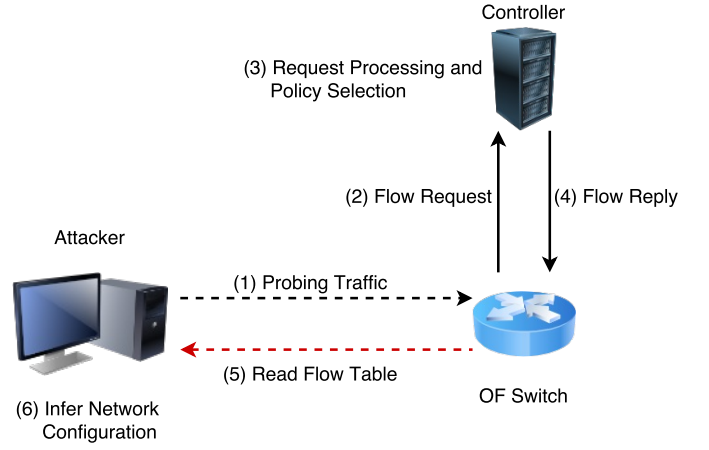


Fig. 2. Overview of a general KYE attack.

The final result is that, through the KYE attack, the attacker is able to learn the control logic of the SDN network regarding a specific type of network flow. Moreover, when the attacker is able to learn the trigger condition for a given network policy (e.g., policy against network scanning attacks), he can then exploit this knowledge to perform additional attacks without triggering detection (see Section VI-B). It is worth noting that the probing traffic might be directed to an end host that is not part of the target network. Indeed, the KYE attack only requires that the probing traffic is routed through a switch of the target SDN. This might be the case, for instance, of an ISP providing defense measures against network attacks to their clients. An attacker could perform a KYE attack to gather intelligence on the ISP network, using as a destination host for the probing traffic the hosts of one of the clients of the ISP. Moreover, introducing probing traffic in the network does not expose the KYE attack to detection. Indeed, one of the strengths of the KYE attack is that it hides behind other attacks: for instance, an attacker performing a KYE attack to gather information related to DoS countermeasures will generate DoS traffic through the target network. When this traffic is detected, the network will simply conclude that a DoS attack is underway, and the KYE attack will remain undetected. Effectively, the probing traffic creates a smokescreen that hides the KYE attack to the network, making it extremely hard to detect in any reliable fashion.

V. KYE INSTANCES

As we discussed before, the KYE attack is a general attack strategy, and the details of the attack vary based on the specific information the attacker wants to gather. In this section, we discuss different instances of the KYE attack, with respect to what type of information the attacker wants to obtain. Section V-A presents the KYE attack in relation to the disclosure of security-related network configuration information. In Section V-B we discuss how the KYE attack can be used to gather intelligence on SDN-related network configuration. Finally, in Section V-C we analyze the use of the KYE attack to disclose general network configuration information.

Additionally, we provide a non-exhaustive list of concrete examples of KYE attack, showing how an attacker can exploit the flow table side-channel to infer different configuration features of the network.

A. Gathering Network Security Configuration Information

When planning an attack, knowing what detection and defense mechanisms are used by the target network is obviously invaluable to an attacker. In this section we discuss how, through repeated probing and analysis of the flow table, an attacker can infer detection mechanisms and defense measures in place in an SDN network for different types of attack. Due to space limitations, in our analysis we focus on some popular SDN-based defense mechanisms proposed in the literature [10], [39].

1) *Worm Infection/Scanning*: Scanning is one of the main preliminary intelligence gathering techniques, used by attackers to gather information about a given target network and by worms to detect vulnerable targets to spread the infection to. Through scanning, an attacker can learn about the number, type and address of hosts in a network, along with what services are offered on which port. This information is a prerequisite for mounting more complex attacks. Therefore, being able to detect and mitigate scanning is extremely important for any network.

KYE Attack. In SDN, an attacker can infer information regarding the type of defense mechanisms used to mitigate scanning and, depending on the detection mechanism employed, the detection threshold for scans. In order to infer such information, an attacker simply needs to send scanning probes from a spoofed address IP_A , varying the characteristics of the scan (e.g., increasing scanning rate, different duration of the scan, various success/failed connection ratios). After each probe, the attacker reads the content of the flow table of the OpenFlow switch and takes note of the new flow rule installed in response to the probe. As long as the scanning is not detected, the flow rules installed simply instruct the switch to forward the traffic coming from IP_A towards different exit ports, based on the destination address. When the scanning rate becomes high enough or the scanning activity lasted long enough for the attack to be detected [35], the controller will install flow rules implementing an appropriate defense measure against scanning.

Detection and Defense Mechanism Inference. Depending on the defense measure used by the network, different types of flow rules will be installed, for instance: traffic filtering [28], rate limiting [35], [46], honeypot redirection [43], or whitehole network approaches [13], [44]. All these defense mechanisms require the installation of very specific flow rules, which differ from the normal flow rules installed when no attack is detected (see Section V-D). Consequently, by recognizing this change in the type of rules installed, the attacker can learn that the network scanning was detected. Moreover, the flow rules required by these mechanisms are easily identifiable and, just by looking at what rule is installed, the attacker

can infer the defense mechanism used by the network (see Section V-D). Finally, for some defense mechanisms that are activated on-demand by the controller, the attacker can also infer the traffic features that trigger detection by the network. For instance, if we consider TRW-CB [38], which is one of the most frequently used anomaly detection algorithms [15] and is already implemented in SDN [35], [43], the attacker can learn the ratio between successful/unsuccessful connections used as detection criteria through repeated probing. Once the attacker discovers the detection threshold, he is then able to carry out the network scan undetected in a second phase, by alternating unsuccessful scanning with successful connections.

Implementation. In Section VI-A we report on the implementation of this instance of the KYE attack. We show that, for realistic detection and defense mechanisms, an attacker is able to learn both the scanning traffic features that trigger detection and the defense mechanism applied in response.

2) *Denial of Service*: DDoS attacks are one of the most widespread type of attacks and their diffusion and sophistication increases every year [5]. Indeed, attackers often use DDoS attacks as a smokescreen to cover data theft and malware installation on target systems [5]. Moreover, the financial damage these attacks cause can range in hundreds of thousands of dollars in peak hours [9]. As a consequence, most organizations adopt one or more DDoS detection and mitigation system [9]. In the context of SDN, DDoS detection schemes tend to employ lightweight mechanisms, such as threshold and entropy-based systems [28], [23] to avoid overloading the controller. Other more complex and computationally expensive approaches do exist, like [18], where the authors employ machine learning techniques to detect possible DoS attacks. Depending on which detection mechanism is used, an attacker can perform a KYE attack to learn if a DoS detection mechanism exists, what defense measure is applied by the network and potentially even the detection criteria employed.

KYE Attack. The KYE Attack for DoS detection is very similar to the one used for network scanning; in the probing phase, the attacker starts a DoS with a low attack rate, simulating a behaviour that is as close as possible to that of a legit client, then gradually increases the profile of the attack. Throughout the attack, the attacker monitors the flow rules installed by the controller on the OpenFlow switch, looking for a change in the flow rules pushed. Indeed, under normal circumstances the controller will simply instruct the switch to route the traffic towards the destination host. However, when the DoS attack is detected, the controller pushes different rules based on the defense mechanism employed by the network, allowing the attacker to learn that the DoS was detected.

Detection and Defense Mechanisms Inference. Defense mechanisms proposed against DoS attacks, like traffic redirection [34], rate limiting [46], [35] or traffic filtering [28], require very specific flow rules that an attacker can easily distinguish from normal routing rules (see Section V-D). Therefore, as soon as these security flow rules are installed on

the monitored OpenFlow switch, the attacker will know that the DoS was detected. By analyzing the specific flow rules installed, he can also infer the defense mechanism applied (see Section V-D). Additionally, when the attack is detected, the attacker can try to infer the detection criteria used by the network. Indeed, for certain type of detection mechanisms such as threshold [23] or entropy-based [28], it is possible to find a good approximation of when exactly an attack is detected. In order to learn these detection thresholds, an attacker can repeat the probing phase several times varying the characteristics of the attack. Upon detection, the attacker will log such characteristics, like duration, attack rate and number of packets sent from each IP address, for instance. After obtaining a sufficiently large sample, the attacker can look for correlations between the characteristics of the detected attacks to learn approximately what values trigger the detection. Even in case of more complex detection systems based on machine learning [18], the attacker can still obtain knowledge about the traffic features used for detection. Indeed, previous work on the area of machine learning shows that it is possible to infer meaningful information about the training set of a classifier [16], which in our case would reveal information about which flows are considered malicious or benign.

3) *Access Control*: Access control mechanisms, like firewalls, are the first and most basic defense mechanism used by networks to enforce security policies [12]. Through its centralized view of the network and distributed enforcement of rules, SDN provides the optimal functionality to implement a consistent distributed firewall in the network [29], [45], [47]. Given that such access control systems are the first defense mechanism that an attacker needs to bypass before attacking a network, learning the exact configuration of such devices would provide a huge advantage in preparing an attack. While this would be extremely challenging in classical networks, due to the fact that access control rules are relegated only to specific security devices, this task is considerably simpler in SDN.

KYE Attack. By performing a KYE attack, the attacker can infer all the access control policies he is interested in by simply probing the monitored OpenFlow switch. In its most basic form, the attacker will send a probe packet to test any given access control policy. For instance, the attacker can try to connect to a protected service using a set of different IPs, in order to understand which subnets are allowed to access that service. For each of these probes the controller will push either a forwarding flow rule, to forward the packets to their destination, or a *drop* rule if the access is not allowed [12]. By repeating the probing for all interesting services and using different source IP addresses, the attacker can map which addresses (or address ranges) are allowed towards/from a certain critical service.

Defense Mechanism Inference. For access control enforcement, the policy that is most used in general is to drop unauthorized traffic flows [12]. If such defense mechanism is in place, the attacker is able to recognize it immediately just

by reading the new flow rule installed (see Section V-D). SDN also allows for more complex defense mechanisms to enforce access control, like traffic redirection towards a honeypot/IDS for instance [34], [43]. As we discuss in Section V-D, the attacker can easily identify even this more complex mechanisms just through observation of the OpenFlow switch flow table.

Implementation. In Section VI-B we report on the implementation of this instance of the KYE attack. We show that, given an access control mechanism, an attacker is able to learn the complete access control matrix enforced by the controller.

B. Gathering SDN-Related Configuration Information

Through a KYE attack, an attacker can infer vital information about SDN-related network configuration, such as flow table management [11], [22] and control plane scalability measures [13], [44]. Knowledge about the configuration of such critical systems provides an attacker with a wider attack surface, as well as allowing him to better focus his resources during an attack.

1) *Flow Table Saturation*: Flow rules, and therefore flow tables, are the core enablers of SDN in OpenFlow. While fine-grained flow rules allow for targeted policy enforcement, the number of flow rules that can be installed on OpenFlow switches is limited. If the flow rule limit is reached (e.g., as a result of a deliberate saturation attack), the switch will not be able to accept rules for new inbound network flows, which will be ignored. To mitigate this vulnerability, the control plane can employ wildcard rules to aggregate the management of multiple network flows with a single flow rule [11], [22]. From the point of view of an attacker, learning under which conditions the controller uses wildcard rules and, more importantly, how to force it to install targeted flow rules, is essential to successfully mount saturation attacks.

The main consideration behind aggregate network flow management is that, in general, it is important to route flows on an individual basis only under certain conditions. In particular, for the purpose of network engineering, this conditions are generally related to the weight of a specific network flow; network flows that surpass a certain threshold (be it a data rate [11] or size [22] threshold), are marked for individual routing through targeted flow rules.

KYE Attack. Through a KYE attack an attacker can infer the thresholds used to classify network flows, allowing him to flood the network with flows for which the controller will generate targeted flow rules. Indeed, similarly to a KYE attack against threshold based detection systems (see Section V-A1), an attacker can easily detect if such aggregation systems are in place and, if so, what thresholds they employ. As a first step, the attacker reads the OpenFlow switch flow table and creates a new network flow not matching any flow rule present. If network flow aggregation is in use, the new flow rule installed on the switch will be a wildcard flow rule, otherwise it will be a flow rule targeted specifically at the new network flow. If a wildcard rule is installed, the attacker can then increase the data rate transmission for that network flow

until a new, flow-specific rule is installed by the controller. At this point, the attacker knows the exact attack rate and/or size of the flow (cumulative amount of data transmitted) required for the controller to install a targeted flow rule.

2) *Control Plane Scalability*: An important issue in SDN lies within the scalability of its control plane. Indeed, previous research demonstrates that, due to the communication required between data and control plane, SDN is subject to a particular type of DoS attack known as control plane saturation attack [13], [44]. This attack aims at overloading the control plane with flow requests by flooding OpenFlow switches with a high number of unique network flows. To mitigate this problem, researchers proposed data plane extension modules aimed at validating the inbound traffic before forwarding a flow request to the controller [13], [44]. For an attacker, detecting if such countermeasures exist is fundamental in order to successfully attack the target network. The core idea behind these countermeasures is to perform the traffic validation at the data plane, before contacting the controller. In [44], for instance, the traffic validation is done through the use of SYN proxy as soon as an OpenFlow switch receives a SYN packet.

KYE Attack. An attacker can detect the presence of this kind of defense mechanisms by means of the KYE attack; indeed when SYN proxy techniques are used, the OpenFlow switch answers to any SYN packet without a flow rule being installed by the controller. An attacker can exploit this behaviour by generating a new connection request (SYN packet) for which there are no flow rules already installed in the flow table. Upon receiving a SYN-ACK response, if there are still no matching flow rules for the generated connection request, then SYN proxy countermeasures are in place in the OpenFlow switches. This behaviour exposes the use of SYN proxy techniques at the data plane level to the attacker, who is then free to attack the network through vulnerabilities of the SYN proxy approach [13].

C. Gathering General Network Configuration Information

Beside security and SDN-related applications and mechanisms, networks have numerous general functionalities and policies ranging from network virtualization [24], [40] to traffic shaping and quality of service (QoS) [26], [27]. In this section we discuss how an attacker can gather knowledge about such functionalities through the KYE attack.

1) *Network Virtualization*: Network virtualization is a key component to enable sharing of physical resources. In this context, the abstraction and programmability offered by SDN provide great advantages in the implementation of such network virtualization techniques. Indeed, there have been several proposals in the literature on how to employ SDN to obtain efficient resource sharing in the context of network virtualization [24], [40]. Although they provide considerable benefits, virtualization and resource sharing also create new problems and vulnerabilities. In fact, sharing physical hardware between multiple tenants can lead to leakage of critical

information [37]. From the point of view of an attacker, learning if some type of network virtualization is applied and, if so, gathering as much information as possible from it, is therefore clearly important. When we consider SDN-related virtualization schemes like FlowN [24], network virtualization is implemented through flow tagging: edge routers tag the traffic with an additional header, which is used by the control plane to map a flow with the routing logic specified by the appropriate tenant.

KYE Attack. Through a KYE attack an attacker can learn, for instance, how network flows are managed by other tenants in the network or which tenants he is sharing a physical machine with. Learning flow management for other tenants is straightforward, since all the flow rules for all tenants are installed in the same set of flow tables. The attacker can just read the various flow rules present in the flow table of the OpenFlow switch, differentiating them based on the header for the specific flow (which indicates the tenant owning that flow).

In order to gather knowledge about which tenants are co-resident on the same physical machine, an attacker just needs to generate probe packets directed to a service run by the tenants he is interested in. If the target tenant is co-resident on the same physical machine as the attacker, the new flow rule installed in response to the probe will instruct the switch to output the traffic to the same port it came from. Once the attacker learns if he is sharing the same physical server with his target, he can then mount additional targeted attacks [37].

D. Correlating Flow Rules and Network Policies

Through the KYE attack, an attacker can infer the exact network-level defense mechanism employed against specific attacks. In this section, we present a non-exhaustive set of defense policies [20] that are used in relation to our examples in Section V-A. Furthermore, we explain how an attacker can correlate a sequence of flow rules obtained during the probing phase to the network policy they implement.

a) *Traffic Filtering*: One of the most basic network-level defense mechanisms is traffic filtering. A traffic filtering policy can be employed to mitigate a large range of attacks, including scanning and DoS attacks [13], [28]. In SDN, traffic filtering is implemented simply by installing a *drop* rule matching the offending network flows on OpenFlow switches. An attacker monitoring the flow table of an OpenFlow switch can easily detect the application of such defense mechanism. Indeed, before the policy is applied, the control plane will push on the OpenFlow switch normal flow rules, instructing the switch to forward the inbound traffic based on the destination address. When the attack is detected and the filtering is applied, the control plane will install only a single drop rule for the all the traffic coming from the attacker's IP address.

b) *Rate Limiting*: Rate limiting is a simple, yet effective defense mechanism to mitigate a wide variety of attacks like scanning and DoS [35], [46]. The most basic rate limiters, the ones assigning a maximum bandwidth to a given aggregate of network flows, are immediately recognizable for an attacker

since they are directly defined in the flow rule matching the aggregate network flows [6]. More complex rate limiting approaches like those implemented in [35], [46] limit the rate of new network flows by delaying the installation of flow rules. In particular, [46] introduces the notion of working set: for each given host, its working set is defined as the set of recently contacted hosts. Whenever a host creates a new network flow addressed at a host outside its working set, the controller will withhold the installation of the corresponding flow rule on the switch for a certain time. After this waiting time expires, the controller instructs the switch to forward the network flow without installing a flow rule. Only when the switch receives a positive reply from the destination host, the controller will install a new flow rule on the switch. An attacker can infer the presence of this defense mechanism by constantly probing the flow table of the switch after creating a new network flow. If rate limiting is in use, the attacker will notice that, even for extremely distant hosts, he will receive a response as soon as the flow rule is installed in the OpenFlow switch. Conversely, when no such technique is used, there will be a delay between the installation of the flow rule on the switch, which would happen as soon as the packet from the attacker reaches the OpenFlow switch, and the moment the reply is received. Therefore, by monitoring the delay in receiving a response after a flow rule is installed, the attacker can infer the use of this defense mechanism.

c) *Whitehole Network*: Another defense mechanism proposed in the literature is SYN proxy [13], [44]. SYN proxy techniques aim at countering the SDN-specific control plane saturation attack [13], [44]. Additionally, SYN proxy implements a whitehole network [44] at the switch level, providing mitigation also against network scans. When such countermeasures are employed, the scenario is slightly different since these are pro-active techniques that are always active, rather than triggered by the installation of a flow rule. Even in this case though, the attacker can infer the existence and the exact type of the defense mechanism employed by the network. Indeed when SYN proxy techniques are used, the attacker will receive a response SYN-ACK packet *without a flow rule being installed on the OpenFlow switch* [44]. Additionally, the attacker will always receive a response SYN-ACK packet to each and every of his probes, even if directed to himself. This behaviour exposes the use of proxy techniques at the data plane level to the attacker, who can attack the OpenFlow switches through vulnerabilities of the proxy approach [13].

d) *Traffic Redirection*: Traffic redirection is a popular defense mechanism since it provides many opportunities to defend a system [34], [41], [43]. For instance, a defender can opportunistically route malicious traffic towards a honeypot, allowing it to isolate the attacker and to study his behaviour [43]. In SDN these defense mechanisms are easy to detect for an attacker. Indeed, by monitoring how the control plane updates flow rule entries for some given network flows, the attacker can infer if his attack traffic is diverted towards a security middlebox/honeypot, nullifying the effect of the countermeasure. In order to do so, an attacker first

generates a new legit network flow towards a given destination D_1 which is routed through a port P_i on the entry switch. The attacker then repeats this step with different destinations, until for a given destination D_n the controller pushes a flow rule instructing the switch to output the matching flow on a port $P_j \neq P_i$. As a second step, the attacker generates probing traffic with a high profile (e.g., high scanning or DoS rate) towards destinations D_1 and D_n for a length of time, observing the flow rules installed in response. If traffic redirection techniques are in place, the control plane will install on the switch a new rule diverting the attack traffic towards the remote middlebox/honeypot. Therefore, all the attack traffic will be routed through the same output port on the OpenFlow switch. Since in the first phase the attacker selected the destinations D_1 and D_n such that packets towards them would be outputted on different ports, if all attack packets towards those same destinations are tunneled through the same output port, then a redirection mechanism is present in the network.

VI. KYE IMPLEMENTATIONS

In order to prove the feasibility and effectiveness of the KYE attack, we implemented two instances of the attack on a test network. In this section, we present the detailed implementation of our attacks, that were aimed at disclosing:

- 1) The presence of a scanning detection and defense mechanism in the network. If present, we also wanted to estimate the detection threshold.
- 2) The presence of a subnetwork access control mechanism. If present, we wanted to learn the subnetwork access control matrix.

Figure 3 depicts the setup used for the evaluation, which includes: the attacker h_0 , a single OpenFlow switch s_1 connected to the controller c , and 100 legit hosts $h_1 - h_{100}$. Hosts $h_1 - h_{100}$ represent known web servers that always reply to connection requests. Hosts $h_1 - h_{100}$, if needed, may be used by the attacker to obtain different connection success ratios during probing, and are not necessarily part of the target network.

In order to simplify the simulation, in our experiments the target network is comprised only of the switch s_1 and the network controller c . Even though in our test network we deployed only a single controller and a single switch, our experiments do not lose generality. This is because the logic used by the controller is the same that would be used in a more complex network, and the rules pushed by the controller are also the same. It is worth noting that it is not a requirement for the attacker to be directly connected to the OpenFlow switch, nor it is for hosts $h_1 - h_{100}$. This is just a simplification we adopted in order to run our simulation. In this test network, we implemented the TRW-CB scanning detection algorithm, which is one of the most used anomaly detection algorithms [15]. Our implementation of TRW-CB follows the SDN implementation detailed in [35]. The test network also implements an IP-based access control mechanism. We ran

all our experiments in a simulated network using the Mininet network simulator [4] and the POX network controller [8].

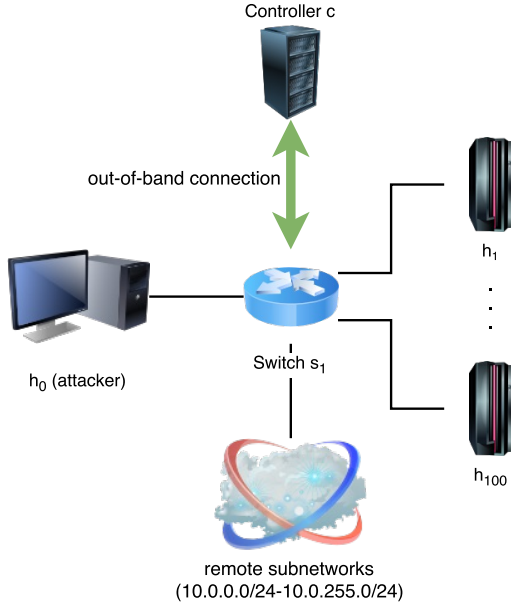


Fig. 3. Experimental setup used for the evaluation.

A. Disclosing Scanning Detection and Defense Mechanisms

The target network implements TRW-CB at the control plane as a scanning detection mechanism, and traffic filtering as a defense measure (see Section V-D). TRW-CB employs both credit limiting, used to limit the amount of first contact connections pending, and sequential hypothesis testing to detect scanning hosts. In particular, we configured the TRW-CB algorithm with the same parameters used in the original paper [38] (base credit of 10, false positive rate ≤ 0.00005 , precision ≥ 0.99).

As an attacker, the first step is to learn if the target network has any kind of scanning detection mechanism in place. To this end, we first initiated a scan with high packet/sec ratio using a spoofed IP, towards a remote subnet which is protected by the target SDN. At the same time, we constantly monitored the flow table for a flow rule implementing a defense mechanism, which would indicate the presence of a scanning detection mechanism. With our configuration, after the first 10 connection attempts failed, TRW-CB correctly identified our probes as scanning activity. Upon detection, the controller pushed a drop flow rule matching all packets coming from the attacker's IP address. Since we were monitoring the flow table of the switch, we detected the rule installation and concluded that the network indeed implements a detection mechanism for scanning attacks, as well as that it uses traffic filtering as a defense measure.

After confirming the existence of a scanning detection mechanism, the following step is to learn which detection criteria are used. As discussed in Section V-A1, in order to do so we initiate several batches of network scans with different

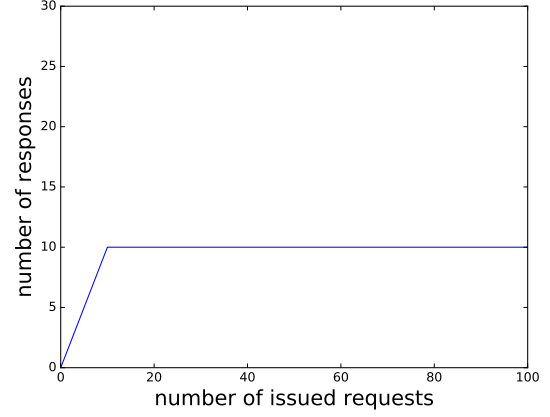


Fig. 4. Number of responses received vs. number of requests issued towards h_1, h_{100} at each batch.

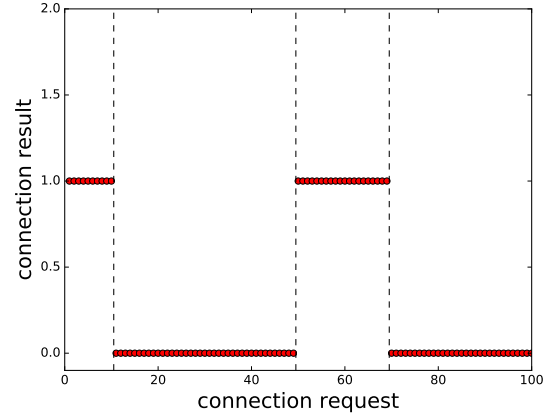


Fig. 5. Connection result for each request of the batch sent towards a known host. 1 indicates a successful connection, 0 a failure (i.e., no response received).

characteristics, like scan rate and successful/failed connections ratio. The results of these batches of scans show two visible characteristics:

- 1) The scanning activity is detected regardless of the scanning rate. This behaviour excludes rate-based scanning detection mechanisms.
- 2) For scan batches where connection requests were sent only towards h_1, h_{100} (which should all send back a response), we received replies only from some of them, as illustrated in Figure 4.

The behaviour shown in Figure 4 is consistent with rate limiting techniques, where connection requests sent at a rate above a certain threshold are dropped. To investigate this anomaly, we started a new scan towards h_1, h_{100} with slightly lower rate. The results are illustrated in Figure 5. As we can see, connections are allowed in bursts: replies were received for the first 10 connections, then 39 requests were dropped, after which connection attempts 50 through 69 were successful, and then connections were dropped once again. Since the scanning

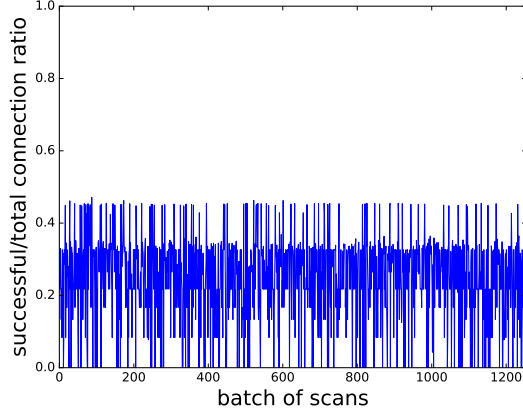


Fig. 6. Ratio of the number of successful connections over the number of total connections, for each batch of scans which resulted in detection.

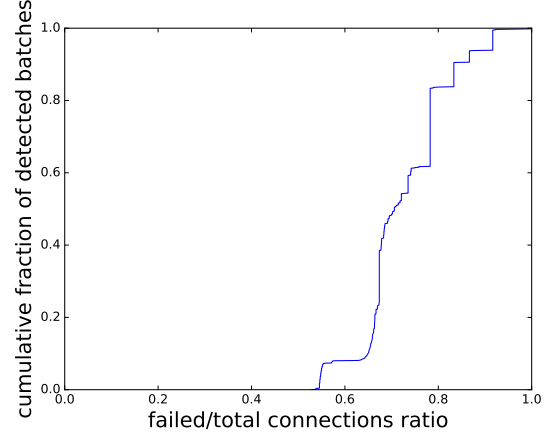


Fig. 7. Empirical cumulative distribution function of the ratio of failed connections over total number of connections, for batches of scans which resulted in detection.

rate was constant for all the 100 connection attempts, this behaviour excludes a standard rate limiting technique. Indeed, from Figure 5 we can see how a host is allowed to contact up to 10 new hosts (i.e., 10 starting credits), after which connections are blocked until pending replies are received. Upon receiving the replies, new credits are allocated to the host, whose connections are correctly forwarded once again. This pattern is consistent with the presence of a credit based rate limiting mechanism [38].

At this point, through the KYE attack, we learned that:

- The network is using a detection mechanism for scanning, which is not rate based.
- The network is using drop rules as a defense mechanism against scanning. Moreover, the network employs an additional preemptive defense measure in the form of credit based rate limiting. Each host is assigned a starting balance of 10 credits and for each successful connection the hosts receives 2 additional credits (the initial 10 successful connections allowed 20 more connections after replies were received).

In the final step of the KYE attack, we initiated several batches of scans, with varying successful/failed connection ratios and scan duration. Each scanning batch terminated either after all planned scans were performed, or abruptly upon detection. Since we are only interested in the characteristics of the scanning attacks that are detected, we isolate detected batches from undetected ones. For the batches of scans that were detected, figure 6 shows the ratio of successful connections over the total number of connections issued and Figure 7 shows the cumulative distribution function of the ratio of failed connections over total number of connections. As these two figures show, the scanning detection criteria employed by the network is clearly based on the ratio of successful and failed connections. Indeed, from Figure 6 we see that network scans are never detected when the ratio of successful connection over the total number of issued connections is above ~ 0.45 . Conversely, from Figure 7 we can see that network scans are

detected when the ratio of failed connections over total number of issued connections is above ~ 0.55 , and never detected when it is below that threshold.

B. Disclosing the Subnetwork Access Control Matrix

After learning the scanning detection criteria used by the network through a KYE attack, we show that we can also infer the complete access control matrix used by the network without being detected by the controller. In our test network, we configured the POX controller with a set of static access control policies, where access to a certain subnetwork is allowed only from a subset of all subnetworks. Whenever a connection request from an unauthorized address is received, the controller instructs the switch to drop the packet without installing any flow rule. If the connection request is from an authorized address, the controller installs a normal forwarding flow rule on the switch for subsequent packets. In this setting, we perform a KYE attack, sending scan probes from each subnetwork to every other. Since with the previous attack we inferred the detection criteria used by the network for scans, *we can now perform this network scan attack completely undetected*.

The attack itself is very simple: at each scan probe, we spoof the source address to make it look like the source of the connection is part of a given subnetwork. We repeat the scan for each pair of source and destination remote subnetworks in the range $10.0.0.0/24 - 10.0.255.0/24$, while opening enough successful connections to remain below the detection threshold. After each scan, we read the flow table of the switch to detect which rule is installed. By monitoring the flow table, we see that no flow rules are installed when the source IP is not allowed to access the subnetwork, while a forwarding rule is installed when the access is authorized. Through this observation we are able to build the access control matrix illustrated in Table I, which reflects exactly the access control rules that were set up at the network controller.

	10.0.0.0/24	10.0.1.0/24	10.0.2.0/24	10.0.3.0/24	10.0.5.0/24	10.0.8.0/24	10.0.10.0/24
10.0.0.0/24	✓	✓	✗	✗	✗	✗	✗
10.0.1.0/24	✓	✓	✗	✗	✗	✗	✗
10.0.2.0/24	✗	✗	✓	✓	✗	✗	✗
10.0.3.0/24	✗	✗	✓	✓	✗	✗	✓
10.0.5.0/24	✗	✗	✗	✗	✓	✓	✗
10.0.8.0/24	✗	✗	✗	✗	✓	✓	✗
10.0.10.0/24	✗	✗	✗	✓	✗	✗	✓

TABLE I

SUBNETWORK ACCESS MATRIX FOR THE TARGET NETWORK, LEARNED BY THE ATTACKER THROUGH A KYE ATTACK. NOTATION ✓: ACCESS ALLOWED; NOTATION ✗: ACCESS RESTRICTED.

VII. COUNTERMEASURE TO THE KYE ATTACK

When considering potential countermeasures, the main problem is that the behaviour responsible for the vulnerability exploited by the KYE attack is also the main strength of SDN: programmability. Therefore, it is not possible to just remove or sensibly alter such behaviour without denaturing SDN itself. Additionally, classical techniques used to avoid eavesdropping, like randomized routing for instance [25], are not applicable against the KYE attack (see Section IX).

In this section, we propose a countermeasure to the KYE attack that does not require modifications to SDN, but rather that takes advantage of SDN programmability. We call this countermeasure *flow obfuscation*. In this section, we use a stronger attack model than the one presented in Section III. While in our presentation of the attack we used the threat model that is least favourable to the attacker (i.e., the attacker has a side-channel only for a single switch), for our countermeasure we consider a threat model that is least favourable to the defender: we assume that the attacker can obtain a flow table side-channel for up to n switches in the target SDN network.

A. Flow Obfuscation

In order to successfully mount a KYE attack, an attacker needs to be able to correlate (i) the network flows he generates to (ii) the reaction they cause in the network. In SDN, this corresponds to the installation of a specific flow rule that the attacker can detect, and from which he can obtain some knowledge about the network. Therefore, if it was possible to prevent the attacker from understanding which network flow caused the installation of which flow rule, the KYE attack would become unfeasible. In order to achieve this goal, we exploit the ability of OpenFlow switches to modify packets in transit. Figure 8 provides an overview of this countermeasure.

Any time a new network flow f_i is received by a switch, the controller installs a *single* flow rule on the switch, with two actions: the first action instructs the switch to modify some header fields of the packets of f_i (e.g., source and destination IP), while the second action tells the switch which output port to use for packet forwarding. This process is repeated for the first $k - 1$ switches s_1, \dots, s_{k-1} in the path, after which, at switch s_k , the controller installs a rule to enforce the appropriate network policy for f_i . The goal of

this countermeasure is to prevent the attacker from learning which network flow causes the installation of a given flow rule; indeed, since the attacker can control up to n switches, when $k > n$ he will never be able to obtain the complete knowledge required for a successful KYE attack. In fact, when $k > n$ there are two possible scenarios:

- If the attacker monitors s_1, \dots, s_{k-1} , then he does not monitor s_k , which is the switch applying the network policy flow rule.
- If the attacker monitors s_k , then he can not know if the installation of a rule on s_k is caused by a network flow he generated. This is because the previous $k - 1$ switches modify the packets at each step, and the attacker does not monitor all of them.

Please note that the use of source and destination IP to identify network flows is just used to present the countermeasure. In general, network flows can be matched on more or even completely different header fields. Even in such cases, the flow obfuscation countermeasure can be implemented by modifying these fields through *set_field* actions [6].

1) *Choosing the Value of K*: While in order to prevent with 100% probability all possible KYE attacks k needs to be greater than n , in the general case this is not required. Indeed, depending on the average out-degree of the switches in the network, even with a value of $k \leq n$ the probability of the attacker obtaining a side-channel on the exact k switches used for flow obfuscation can be low. Given an SDN network, let us assume that $o + 1$ is the average out-degree of a switch and that an attacker knows the number of switches used in flow obfuscation. If an attacker a can monitor at most n switches (for simplicity, we assume $n \bmod k = 0$), the probability P_{s_1, s_k}^s of him monitoring exactly the k switches s_1, \dots, s_k used for flow obfuscation is:

$$P_{s_1, s_k}^s = \left(\frac{(n/k)}{o} \right)^{k-1},$$

since attacker always knows the identity of s_1 (the edge switch through which the attacker's traffic is routed). A network manager can decide the appropriate value for k based on the expected value of n and the maximum probability P_{s_1, s_k}^s that he is willing to accept.

2) *Limitations*: The flow obfuscation countermeasure can effectively prevent the KYE attack and is configurable to the

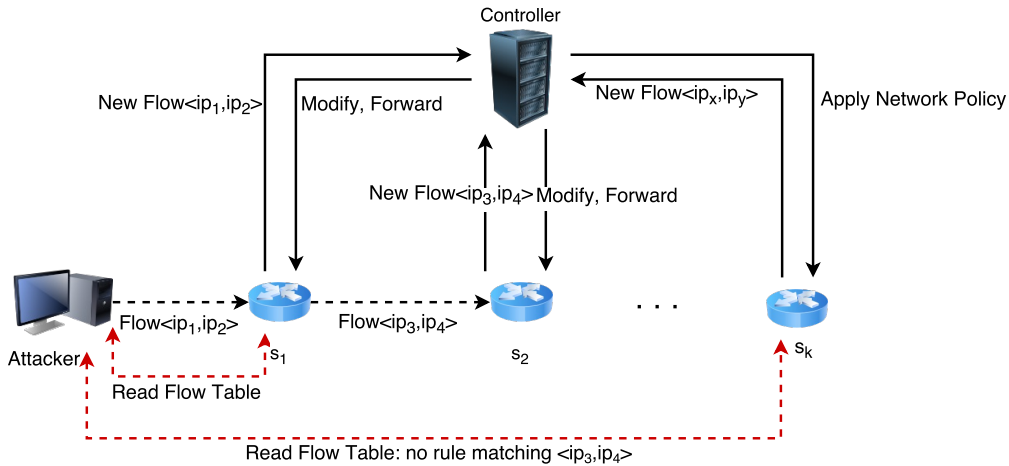


Fig. 8. Overview of the flow obfuscation countermeasure. The red dashed lines indicate the flow table side-channel.

needs of the network, but also presents some drawbacks. The first drawback is related to the delay in the application of network policies to network flows. Indeed, while in a normal SDN environment network policies are applied immediately at the ingress switch, when flow obfuscation is in place the network policy is delayed to the k -th hop in the network. While this behaviour might not impact all flows, in case of flows that require less than k hops to reach the destination (e.g., flows that would be immediately dropped) it adds additional load on the network. The second drawback is related to the additional work required at the control plane. In fact, in order to correctly route the network flows the controller needs to keep track of each header modification applied by the switches. The higher k is, the higher the amount of bookkeeping the controller has to do to correctly identify the network flows.

These drawbacks can be mitigated by applying flow obfuscation in an intelligent manner. Indeed, the length of the flow obfuscation path can be tailored to the average length of network flows in the network; this way, on average flow obfuscation will add only a negligible delay to the application of network policies and the controller will incur a reduced overhead. Moreover, while on average flow obfuscation will indeed increase the load on network links, this load can be effectively distributed by the controller by dynamically changing the flow obfuscation path. This allows the controller to avoid link saturation, even in case of deliberate DoS attacks, thanks to the global network visibility provided by SDN. Additionally, it is worth noting that when applying flow obfuscation, the controller incurs overhead only for the first packet of each network flow. In fact, once the header-modification and forwarding flow rules are installed on the switches, all packets will be modified and correctly routed without the intervention of the controller.

VIII. LIMITATIONS OF THE KYE ATTACK AND FUTURE WORK

The KYE attack requires constant assessments of the flow table by the attacker and the ability to recognise deviations in

the type of flow rules installed. A human attacker can find this process bothersome and non-trivial in certain cases, leading to imprecise identification of detection conditions and/or defense measures applied. Additionally, the attacker might not be able to correctly identify the network policy applied if the policy itself is new or unknown to the community. While this might be the case, we argue that the attacker does not necessarily need to learn the exact nature of the network policy applied, as long as he is able to learn enough information for his purposes. For instance in case of network scanning, not being able to identify the defense mechanism used might be acceptable for the attacker, as long as he is able to understand when the scanning traffic is detected and when it is not. The attacker can learn when the scanning is detected by observing a deviation in the flow rules installed for different types of probing traffic (e.g., very low scan rate probing traffic v.s. fast scan rate probing traffic).

In our future work, we plan on automating the KYE attack by means of machine learning techniques. Indeed, the efficacy of the KYE attack is based on the ability to recognising patterns in the features of the generated traffic and in the flow rules installed in response, as well as to recognise significant deviations from such patterns. We believe that machine learning can be successfully applied to detect the baseline behaviour of the control plane under normal traffic conditions. Once this baseline behaviour has been identified, we can use a classifier to detect and categorize deviations from such behaviour triggered by specific traffic flows. It would then be possible to use another classifier to find the most relevant features of the attack traffic which caused the abnormal network state (e.g., detection of an attack). In this system model, the OpenFlow switch acts as an oracle: the switch can be repeatedly queried to learn if some network flows trigger a deviation in the usual type of flow rules installed. By analyzing the difference between the features of network flows that triggered the abnormal network state versus those which did not, it is possible to learn which are the characteristics that caused the reaction.

In complex real-world networks, it is possible to have several subnetworks, each with separate network policies. In this scenario, the KYE attack allows to learn the network configuration of the subnetwork for which the attacker has a flow table side-channel. If the attacker wishes to learn the configuration of other subnetworks, he will need to obtain a flow table side-channel for one switch on each of those subnetworks.

IX. RELATED WORK

SDN has become a popular research topic in recent years, especially in relation to security. Despite significant research efforts, to the best of our knowledge this is the first work to analyze the vulnerabilities caused by the distributed enforcement of rules in SDN. In [33], Kreutz et al. discuss the effects of an attacker compromising a switch in the network, which can result in traffic injection attacks, man-in-the-middle attacks and traffic filtering. Similarly, in [14] the authors consider an attacker with full control over the switch and discuss several possible attacks like man-in-the-middle, state, and topology spoofing. However, both of these works assume a stronger attacker, with full control over the switch, differently from us. Moreover, in these works [14], [33] the attacker actively modifies the state of the OpenFlow switch, exposing the attack to detection by the controller through querying [17], [31] the state of the switch or RTT analysis [14]. In [32], Klöti et al. analyze the overall security of OpenFlow using the STRIDE methodology. Between other threats, the authors also identify the risk of information disclosure in OpenFlow-based SDN. Indeed, they prove that by analyzing the RTT for a specific network flow, an attacker is able to infer if a flow rule for a specific network flow is already installed in an OpenFlow switch. While the goal of this attack is to obtain some information on the state of a switch, the KYE attack allows to learn a much greater amount of information about the logic and policies of the network, such as attack detection threshold and defense mechanisms applied. Another set of works related to ours pertain to SDN fingerprinting [42]. SDN fingerprinting techniques use RTT to infer if a given network is an SDN or a classical network. These techniques apply the observation, made in [32], that in SDN the first packet belonging to a network flow has a higher RTT than subsequent ones. By exploiting this asymmetry, an attacker can successfully infer if a network is an SDN with high accuracy. Fingerprinting attacks are related to our work in the sense that they also try to gather intelligence about the configuration of a network. However, the approach and the type of information obtained are completely different from our proposal and the amount of information retrieved with the KYE attack is much greater.

The family of countermeasures that may seem applicable to the KYE attack are route randomization countermeasures [25]. These countermeasures aim to provide resistance to reconnaissance and eavesdropping. However, there are two major differences between the KYE attack and a classical eavesdropping scenario that make these countermeasures not applicable.

First, with the KYE attack, the attacker only needs to have a flow table side-channel on the switch through which the attacker's traffic enters the network, rendering randomized routing ineffective. This is because, contrary to a classical eavesdropping scenario where the goal is to sniff the traffic of a victim, the attacker knows which switch in the network will be the gateway for his own traffic. Second, the duration of the network flow in case of the KYE attack and in case of classical eavesdropping are different. Indeed, while in a classical eavesdropping scenario the duration of the network flow is limited, in case of the KYE attack it can last as long as the attacker needs it to. Therefore even if the attacker's traffic is routed through the entry switch only a fraction of the time, this simply means the attacker will have to increase the duration of the probing phase of the KYE attack. In particular, the attacker can opportunistically suspend the probing when the traffic is not routed through the entry switch anymore, and resume when the traffic is routed through it once again. Finally, the KYE attack can be used to learn the characteristics of the randomized routing mechanism itself, like the duration of the randomization intervals.

X. CONCLUSIONS

In this paper, we proposed a thorough analysis of the vulnerability introduced by the on-demand installation of flow rules in SDN. We presented the novel KYE attack which, with minimal requirements, allows an adversary to gather an extensive amount of information regarding the configuration of the network, ranging from security-related aspects to network engineering policies. We implemented the KYE attack and conducted a thorough evaluation, showing its feasibility against a popular scanning detection algorithm and against standard access control policies. Finally, we proposed the flow obfuscation countermeasure to the KYE attack, which provides provable security guarantees and can be tailored to the needs of a specific network under consideration.

REFERENCES

- [1] Hp procurve 5400 openflow switch. http://archive.openflow.org/wk/index.php/Configuring_HP_Procurve. Accessed: 08-2016.
- [2] Juniper screenos authentication backdoor. <https://community.rapid7.com/community/infosec/blog/2015/12/20/cve-2015-7755-juniper-screenos-authentication-backdoor>. Accessed: 08-2016.
- [3] Microsoft's DEMON. https://sharkfest.wireshark.org/sharkfest.12/presentations/A-4_Leveraging_Openflow_to_create_a_Large_Scale_and_Cost_Effective_Packet_Capture_Network.pdf. Accessed: 07-2016.
- [4] Mininet network simulator. <http://mininet.org/>. Accessed: 05-2016.
- [5] Neustar annual ddos attacks and impact report. <https://www.neustar.biz/resources/whitepapers/ddos-protection/2014-annual-ddos-attacks-and-impact-report.pdf>. Accessed: 03-2016.
- [6] Openflow specification. <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-switch-v1.5.1.pdf>. Accessed: 03-2016.
- [7] Openflow wiki. http://archive.openflow.org/wk/index.php/OpenFlow_Tutorial. Accessed: 08-2016.
- [8] Pox network controller. <https://github.com/noxrepo/pox>. Accessed: 05-2016.
- [9] US DDoS attacks and protection report. <https://www.neustar.biz/ddos-attacks-report>. Accessed: 03-2016.

- [10] I. Ahmad, S. Namal, M. Ylianttila, and A. Gurtov. Security in software defined networks: A survey. *IEEE Communications Surveys Tutorials*, 2015.
- [11] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat. Hedera: Dynamic flow scheduling for data center networks. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, NSDI'10, 2010.
- [12] E. S. Al-Shaer and H. H. Hamed. Modeling and management of firewall policies. *IEEE Transactions on Network and Service Management*, 2004.
- [13] M. Ambrosin, M. Conti, F. De Gaspari, and R. Poovendran. Lineswitch: Efficiently managing switch flow in software-defined networking while effectively tackling dos attacks. In *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security*, ASIACCS '15, 2015.
- [14] M. Antikainen, T. Aura, and M. Särelä. Spook in your network: Attacking an sdn with a compromised openflow switch. In *Secure IT Systems: 19th Nordic Conference, NordSec 2014*, 2014.
- [15] A. B. Ashfaq, M. J. Robert, A. Mumtaz, M. Q. Ali, A. Sajjad, and S. A. Khayam. A comparative evaluation of anomaly detectors under portscan attacks. In *Proceedings of the 11th International Symposium on Recent Advances in Intrusion Detection*, RAID, 2008.
- [16] G. Ateniese, L. V. Mancini, A. Spognardi, A. Villani, D. Vitali, and G. Felici. Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *International Journal of Security and Networks*, 2015.
- [17] K. Benton, L. J. Camp, and C. Small. Openflow vulnerability assessment. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, HotSDN '13, 2013.
- [18] R. Braga, E. Mota, and A. Passito. Lightweight ddos flooding attack detection using nox/openflow. In *Proceedings of the 35th IEEE Conference on Local Computer Networks*, LCN, 2010.
- [19] B. D. Carrier and J. Grand. A hardware-based memory acquisition procedure for digital investigations. *Digit. Investig.*, 2004.
- [20] C. J. Chung, P. Khatkar, T. Xing, J. Lee, and D. Huang. Nice: Network intrusion detection and countermeasure selection in virtual network systems. *IEEE Transactions on Dependable and Secure Computing*, 2013.
- [21] H. Cui, G. O. Karame, F. Klaedtke, and R. Bifulco. On the fingerprinting of software-defined networks. *IEEE Transactions on Information Forensics and Security*, 2016.
- [22] A. R. Curtis, J. C. Mogul, J. Tourrilhes, P. Yalagandula, P. Sharma, and S. Banerjee. Devoflow: Scaling flow management for high-performance networks. In *Proceedings of the ACM SIGCOMM Conference*, SIGCOMM '11, 2011.
- [23] M. Dhawan, R. Poddar, K. Mahajan, and V. Mann. Sphinx: Detecting security attacks in software-defined networks. In *Proceedings of the 2015 Network and Distributed System Security Symposium*, NDSS 2015.
- [24] D. Drutskey, E. Keller, and J. Rexford. Scalable network virtualization in software-defined networks. *IEEE Internet Computing*, 2013.
- [25] Q. Duan, E. Al-Shaer, and H. Jafarian. Efficient random route mutation considering flow and network constraints. In *Proceedings of the IEEE conference on Communications and Network Security*, CNS, 2013.
- [26] H. E. Egilmez, S. T. Dane, K. T. Bagci, and A. M. Tekalp. Openqos: An openflow controller design for multimedia delivery with end-to-end quality of service over software-defined networks. In *Proceedings of the Signal Information Processing Association Annual Summit and Conference*, APSIPA ASC, 2012.
- [27] H. E. Egilmez, B. Gorkemli, A. M. Tekalp, and S. Civanlar. Scalable video streaming over openflow networks: An optimization framework for qos routing. In *Proceedings of the 18th IEEE International Conference on Image Processing*, 2011.
- [28] K. Giotis, C. Argyropoulos, G. Androulidakis, D. Kalogeras, and V. Maglaris. Combining openflow and sflow for an effective and scalable anomaly detection and mitigation mechanism on sdn environments. *Computer Networks*, 2016.
- [29] H. Hu, W. Han, G.-J. Ahn, and Z. Zhao. Flowguard: Building robust firewalls for software-defined networks. In *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, HotSDN '14, 2014.
- [30] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat. B4: Experience with a globally-deployed software defined wan. *SIGCOMM Comput. Commun. Rev.*, 2013.
- [31] A. Kamisiński and C. Fung. Flowmon: Detecting malicious switches in software-defined networks. In *Proceedings of the 2015 Workshop on Automated Decision Making for Active Cyber Defense*, SafeConfig '15, 2015.
- [32] R. Klöti, V. Kotronis, and P. Smith. Openflow: A security analysis. In *Proceedings of the 21st IEEE International Conference on Network Protocols*, ICNP, 2013.
- [33] D. Kreutz, F. M. Ramos, and P. Verissimo. Towards secure and dependable software-defined networks. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, HotSDN '13, 2013.
- [34] A. Mahimkar, J. Dange, V. Shmatikov, H. Vin, and Y. Zhang. Dfence: Transparent network-based denial of service mitigation. In *Proceedings of the 4th USENIX Conference on Networked Systems Design and Implementation*, NSDI'07, 2007.
- [35] S. A. Mehdi, J. Khalid, and S. A. Khayam. Revisiting traffic anomaly detection using software defined networking. In *Proceedings of the 14th International Symposium on Recent Advances in Intrusion Detection*, RAID, 2011.
- [36] G. Pickett. Abusing software defined networks. BlackHat, 2014.
- [37] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09, 2009.
- [38] S. E. Schechter, J. Jung, and A. W. Berger. Fast detection of scanning worm infections. In *Proceedings of the 7th International Symposium on Recent Advances in Intrusion Detection*, RAID, 2004.
- [39] S. Scott-Hayward, S. Natarajan, and S. Sezer. A survey of security in software defined networks. *IEEE Communications Surveys Tutorials*, 2016.
- [40] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar. Can the production network be the testbed? In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation*, OSDI'10, 2010.
- [41] S. Shin and G. Gu. Cloudwatcher: Network security monitoring using openflow in dynamic cloud networks (or: How to provide security monitoring as a service in clouds?). In *Proceedings of the 20th IEEE International Conference on Network Protocols*, ICNP, 2012.
- [42] S. Shin and G. Gu. Attacking software-defined networks: A first feasibility study. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, HotSDN '13, 2013.
- [43] S. Shin, P. Porras, V. Yegneswaran, M. Fong, G. Gu, and M. Tyson. Fresco: Modular composable security services for software-defined networks. In *Proceedings of the 20th Annual Network and Distributed System Security Symposium*, NDSS 2013.
- [44] S. Shin, V. Yegneswaran, P. Porras, and G. Gu. Avant-guard: Scalable and vigilant switch flow management in software-defined networks. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security*, CCS '13, 2013.
- [45] M. Suh, S. H. Park, B. Lee, and S. Yang. Building firewall over the software-defined network controller. In *Proceedings of the 16th International Conference on Advanced Communication Technology*, 2014.
- [46] J. Twycross and M. M. Williamson. Implementing and testing a virus throttle. In *Proceedings of the 12th Conference on USENIX Security Symposium - Volume 12*, SSYM'03, 2003.
- [47] J. Wang, Y. Wang, H. Hu, Q. Sun, H. Shi, and L. Zeng. Towards a security-enhanced firewall application for openflow networks. In *Proceedings of the 5th International Symposium on Cyberspace Safety and Security*, CSS'13, 2013.